

Self-Adaptation of Mutation Rates and Dynamic Fitness

Matthew R. Glickman* and Katia P. Sycara

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3891, USA
glickman+@cs.cmu.edu, katia@ri.cmu.edu

Abstract

In any search via artificial evolution, the likelihood of stagnation at local optima is determined by the particular choices of representation and search operators. Because the performance impact of these design choices is difficult to predict, it is an attractive option to let the representation and/or operators themselves evolve. However, effective evolution at this meta-level can be difficult to achieve for a number of reasons, including: (1) The complexity of the search space is increased; and (2) selection acts at the level of the fitness function and only indirectly at the meta-level, favoring variations only to the extent to which they are stochastically associated with fitness improvements. The question then becomes: Under what conditions is evolution of the representation and/or operators likely to be most effective?

To explore this question empirically, we've chosen to search for vectors of weights in an artificial neural network that minimize the sum-of-squared error (SSE) between the computed function and some simple target functions evaluated over a set of example inputs. As an instantiation of evolvable search operators, we've chosen to associate a *variable* mutation rate (*e.g.* as in Bäck 1991) with each weight, providing selection with the capacity to focus variation on specific subsets of the complement of weights.

While our results have quite generally favored the self-adaptation of per-weight mutation rates in the chosen domain, various experiments have indicated that they are most helpful at the point when only a small subset of the vector is ripe for evolutionary tuning. At the same time, it has been found that the variable rate scheme sometimes enhances the likelihood of the population stagnating at a premature local optimum. This effect appears to arise when the population finds its way to a local optimum steep and broad enough that the *expected fitness* of offspring will be greater the more strongly they resemble their parents. At this point, selection will tend to favor those individuals with the most conservative mutation rates, *i.e.*

driving mutation rates toward zero, nestling the population snugly into the local optimum. These observations have led to the conjecture that meta-evolution may perform best when coupled with a dynamic fitness function that differentially favors reduced error in those regions of the target function where the population is performing worst (*e.g.* as in Rosin and Belew 1995). Such fitness functions might both exploit the ability to focus variation on an appropriate subset of the vector, as well as provide enough selective pressure to pop populations out of premature local optima.

Instead of minimizing the SSE, experiments have been conducted in which each individual is assigned a fitness score equal to the sum of its performance on each example point in the function divided by the population's total performance on the same example. Where f_i is the fitness of an individual i , e_{ij} is the error of individual i on example j (maximum value is 1), and with a population of size n :

$$f_i = \sum_j \frac{1 - e_{ij}^2}{n - \sum_k e_{kj}^2}$$

Initial results indicate that dynamism in the fitness function is an advantage, but not in the latter stages of the search, where it tends to inhibit fine-tuning. The best results have been achieved by starting search with the dynamic function, but switching to the static function once the average SSE drops below some pre-set threshold.

References

- Bäck, T. 1991. Self-adaptation in genetic algorithms. In *Proceedings of the First European Conference on Artificial Life*. MIT Press, 1992.
- Rosin, C., and Belew, R. 1995. Methods for Competitive Co-Evolution: Finding Opponents Worth Beating. In Larry J. Eshelman (Ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms*, pp. 373-380. Morgan Kaufmann, 1995.

*Supported in part by an NSF Graduate Fellowship Award